

## Hive 1.0.0发布

美国时间2015年2月4日，Hive 1.0.0正式发布了。该版本是Apache Hive九年来工作的认可，并且开发者们正在继续开发。Apache Hive 1.0.0版本本来是要命名为Hive 0.14.1的，但是社区感觉是时候以1.x.y结构来命名。

虽然被叫做1.0.0版本，但是其中的改变范围很少，主要有两个改变：1、开始为HiveMetaStoreClient定义公开的API(HIVE-3280)；2、HiveServer 1被移除(HIVE-6977)。API文档正在开始，并且会在HIVE-9363里面继续。

HiveServer 1的移除是Hive向企业版准备的重要一步。HiveServer 2在两年半年前由Cloudera贡献给Hive社区的，并且使得Hive支持JDBC、ODBC和细粒度的权限控制。

对于CLI用户来说，迁移到HiveServer 2将需要使用Beeline。Cloudera和Intel将继续投入巨资在Beeline和HiveServer 2上，来使得这个过渡变得更加简单。

下一个Hive主要版本0.15.0将会命名为1.1.0。我们最想要的Hive-on-Apache Spark将会在Hive 1.1.0中发布，这个版本正在开发中。



如果想及时了解Spark、Hadoop或者Hbase相关的文章，欢迎关注微信公共帐号：iteblog\_hadoop

Hive 1.0.0版本主要修改的Bug和提升请参见下面列表：

```
/**  
 * User: 过往记忆  
 * Date: 15-02-06  
 * Time: 上午12:30  
 * bolg:  
 * 本文地址 : /archives/1265  
 * 过往记忆博客，专注于hadoop、hive、spark、shark、flume的技术博客，大量的干货  
 * 过往记忆博客微信公共帐号 : iteblog_hadoop  
 */
```

Bug

[HIVE-5631] - Index creation on a skew table fails

[HIVE-5664] - Drop cascade database fails when the db has any tables with indexes

[HIVE-6198] - ORC file and struct column names are case sensitive

[HIVE-6468] - HS2 & Metastore using SASL out of memory error when curl sends a get request

[HIVE-7270] - SerDe Properties are not considered by show create table Command

[HIVE-8099] - IN operator for partition column fails when the partition column type is DATE

[HIVE-8295] - Add batch retrieve partition objects for metastore direct sql

[HIVE-8374] - schematool fails on Postgres versions < 9.2

[HIVE-8485] - HMS on Oracle incompatibility

[HIVE-8706] - Table statistic collection on counter failed due to table name character case.

[HIVE-8715] - Hive 14 upgrade scripts can fail for statistics if database was created using auto-create

[HIVE-8739] - handle Derby and Oracle errors with joins and filters in Direct SQL in a invalid-DB-specific path

[HIVE-8784] - Querying partition does not work with JDO enabled against PostgreSQL

[HIVE-8805] - CBO skipped due to SemanticException: Line 0:-1 Both left and right aliases encountered in JOIN 'avg\_cs\_ext\_discount\_amt'

[HIVE-8807] - Obsolete default values in webhcat-default.xml

[HIVE-8811] - Dynamic partition pruning can result in NPE during query compilation

[HIVE-8827] - Remove SSLv2Hello from list of disabled protocols

[HIVE-8830] - hcatalog process don't exit because of non daemon thread

[HIVE-8845] - Switch to Tez 0.5.2

[HIVE-8866] - Vectorization on partitioned table throws ArrayIndexOutOfBoundsException when partitions are not of same #of columns

[HIVE-8870] - errors when selecting a struct field within an array from ORC based tables

[HIVE-8873] - Switch to calcite 0.9.2

[HIVE-8876] - incorrect upgrade script for Oracle (13->14)

[HIVE-8880] - non-synchronized access to split list in OrcInputFormat

[HIVE-8886] - Some Vectorized String CONCAT expressions result in runtime error

Vectorization: Unsupported vector output type: StringGroup

[HIVE-8888] - Mapjoin with LateralViewJoin generates wrong plan in Tez

[HIVE-8947] - HIVE-8876 also affects Postgres < 9.2

[HIVE-8966] - Delta files created by hive hcatalog streaming cannot be compacted

[HIVE-9003] - Vectorized IF expr broken for the scalar and scalar case

[HIVE-9025] - join38.q (without map join) produces incorrect result when testing with multiple reducers

[HIVE-9038] - Join tests fail on Tez

[HIVE-9051] - TezJobMonitor in-place updates logs too often to logfile

[HIVE-9053] - select constant in union all followed by group by gives wrong result

[HIVE-9067] - OrcFileMergeOperator may create merge file that does not match properties of input files

[HIVE-9090] - Rename "Tez File Merge Work" to smaller name

[HIVE-9108] - Fix for HIVE-8735 is incorrect (stats with long paths)

[HIVE-9111] - Potential NPE in OrcStruct for list and map types

[HIVE-9112] - Query may generate different results depending on the number of reducers

[HIVE-9114] - union all query in cbo test has undefined ordering

[HIVE-9126] - Backport HIVE-8827 (Remove SSLv2Hello from list of disabled protocols) to 0.14 branch

[HIVE-9141] - HiveOnTez: mix of union all, distinct, group by generates error

[HIVE-9155] - HIVE\_LOCKS uses int instead of bigint hive-txn-schema-0.14.0.mssql.sql

[HIVE-9162] - stats19 test is environment-dependant

[HIVE-9166] - Place an upper bound for SARG CNF conversion

[HIVE-9168] - Vectorized Coalesce for strings is broken

[HIVE-9205] - Change default tez install directory to use /tmp instead of /user and create the directory if it does not exist

[HIVE-9234] - HiveServer2 leaks FileSystem objects in FileSystem.CACHE

[HIVE-9249] - java.lang.ClassCastException:

org.apache.hadoop.hive.serde2.io.HiveVarcharWritable cannot be cast to org.apache.hadoop.hive.common.type.HiveVarchar when joining tables

[HIVE-9278] - Cached expression feature broken in one case

[HIVE-9317] - move Microsoft copyright to NOTICE file

[HIVE-9359] - Export of a large table causes OOM in Metastore and Client

[HIVE-9361] - Intermittent NPE in SessionHiveMetaStoreClient.alterTempTable

[HIVE-9390] - Enhance retry logic wrt DB access in TxnHandler

[HIVE-9401] - Backport: Fastpath for limited fetches from unpartitioned tables

[HIVE-9404] - NPE in

org.apache.hadoop.hive.metastore.txn.TxnHandler.determineDatabaseProduct()

[HIVE-9436] - RetryingMetaStoreClient does not retry JDOExceptions

[HIVE-9473] - sql std auth should disallow built-in udfs that allow any java methods to be called

[HIVE-9514] - schematool is broken in hive 1.0.0

## Improvement

[HIVE-3280] - Make HiveMetaStoreClient a public API

[HIVE-8933] - Check release builds for SNAPSHOT dependencies

## Task

[HIVE-6977] - Delete Hiveserver1

本博客文章除特别声明，全部都是原创！

原创文章版权归过往记忆大数据（过往记忆）所有，未经许可不得转载。

本文链接: 【】()